L1-regularized Reconstruction Error as Alpha Matte

Jubin Johnson, Hisham Cholakkal, and Deepu Rajan

arXiv:1702.02744v1 [cs.CV] 9 Feb 2017

***Abstract*—Sampling-based alpha matting methods have traditionally followed the compositing equation to estimate the alpha value at a pixel from a pair of foreground (F) and background (B) samples. The (F,B) pair that produces the least reconstruction error is selected, followed by alpha estimation. The significance of that residual error has been left unexamined. In this letter, we propose a video matting algorithm that uses L1-regularized reconstruction error of F and B samples as a measure of the alpha matte. A multi-frame non-local means framework using coherency sensitive hashing is utilized to ensure temporal coherency in the video mattes. Qualitative and quantitative evaluations on a dataset exclusively for video matting demonstrate the effectiveness of the proposed matting algorithm.**

***Index Terms*—Residual error, video matting, non-local means.**

I. Introduction

DIGITAL matting refers to the problem of accurate foreground extraction and finds its use in image and video editing. Mathematically, any pixel color I_i can be modeled as a convex combination of the foreground color (F_i) and the background color (B_i) such that

$$I_i = \alpha_i F_i + (1 - \alpha_i) B_i, \tag{1}$$

where α_i is the opacity (alpha) value at pixel i. Determining α is an under-constrained problem, made tractable by means of user-input labels in the form of a trimap or scribbles.

Matting methods are generally divided into sampling-based [1], [2], [3], [4] and propagation-based [5], [6] approaches. The former category uses color values from the known foreground and background regions to find the best foreground-background (F, B) pair that represents the true foreground and background colors to estimate α of a given pixel. Different sampling strategies (local/global) and optimization criteria for selecting the best (F, B) pair distinguish these approaches. Similar color distribution among the foreground and background regions is a challenge since the samples cannot discriminate between F and B regions anymore. Propagation-based methods leverage the correlation between neighboring pixels with respect to local image statistics to interpolate the known alpha values to the unknown regions. As with sampling approaches, false correlations between neighboring F and B pixels occurs due to color similarity. Moreover, strong edges and textured regions fail to propagate the alpha accurately. Recently, deep learning based approaches [7] have shown to perform well in natural image matting.

J. Johnson, H. Cholakkal, and D. Rajan are with the Multimedia Lab, School of Computer Science and Engineering, Nanyang Technological University, Singapore, 639798 (e-mail: {jubin001, hisham002, asdrajan}@ntu.edu.sg).

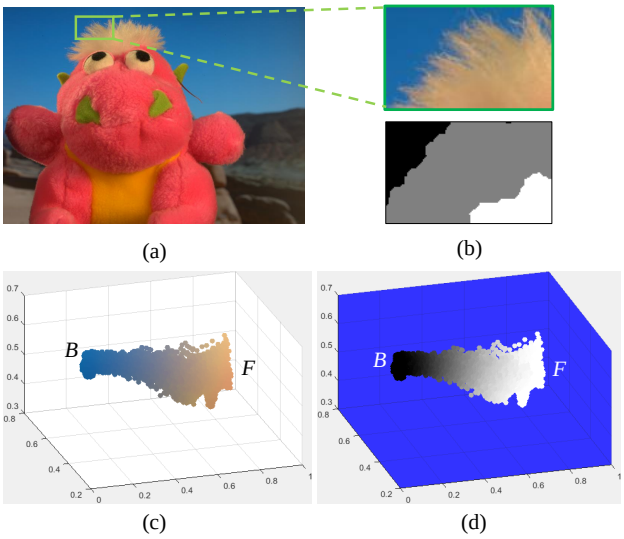

Fig. 1. Illustration of reconstruction error as alpha matte. (a) Input image with a (b) zoomed window and its trimap, (c) color distribution of the pixels in RGB color space indicating definite and unknown regions, and (d) plot of alpha generated from reconstruction error of pixels with values in $[0, 1]$.

Video matting, apart from extracting spatially accurate mattes on each frame, also has the additional requirement of temporal coherence across the video [8], [9], [10]. The human visual system is highly sensitive to jitter and temporal inconsistencies across frames. Low contrast and fast motion are factors that contribute to inaccurate matte in a frame, thereby leading to temporal jitter across the extracted video matte. Although the quality of the mattes obtained by independently applying image matting algorithms to each frame is high, it does not result in temporally coherent mattes. α-propagation has been extended to the temporal domain as post-processing to alleviate this problem. Snapcut [11] uses the matting Laplacian [5] to bias the alpha to the previous frame. A motion-aware Laplacian is constructed to propagate the matte temporally in [12]. Level-set interpolation is used to temporally smooth the estimated mattes in [9]. Optical flow is used to warp the alpha from the previous frame in the Laplacian formulation in [8].

The proposed approach is based on sampling. As mentioned earlier, sampling methods find the best (F, B) pair that satisfies eq. (1) and use it to estimate the alpha value. The reconstruction error in the selected pair is $\xi_i = \|I_i - (\hat{\alpha}_i F_i + (1 - \hat{\alpha}_i) B_i)\|$. The significance of this residual error for matting has largely been left unexamined in literature. Johnson *et al.* [10] showed sparse coding as an alternative to

(a) (b) (c) (d)

Fig. 2. L1-regularized reconstruction error as alpha matte. (a) Input image and its trimap, plots of (b) ξ_F, (c) ξ_B and (d) estimated alpha map using eq. (5).

the compositing equation for estimating the α value at a pixel. Inspired by this, we propose a sampling-based approach that looks at matting from the perspective of sparse reconstruction error of feature samples. Fig. 1 illustrates the motivation behind using reconstruction error as a measure of the matte in a real image. A zoomed region of the input image in Fig. 1(a) and its trimap are shown in Fig. 1(b), representing a hairy region containing mixed pixels. The local smoothness assumption between the alpha values of neighboring pixels is paramount to extracting a good matte. In a real image, alpha would gradually transition between the definite F and B with the true mixed pixel alphas having an intermediate value. The RGB color distribution of pixels in the image patch varies smoothly between the foreground and background with the blending peaking at the middle of the unknown region (Fig. 1(c)). Similarly, the error obtained during reconstruction using F and B samples can be thought of as a probability measure that varies smoothly between the foreground and background regions, gradually rising from the definite regions and peaking at true mixed pixels. As can be seen in Fig. 1(c) and (d), the color distribution of pixels in a real image and the residual error are highly correlated. To the best of our knowledge, we are the first to formulate matting from the perspective of reconstruction error. A patch-based non-local means (NLM) framework using coherency sensitive hashing across multiple frames is integrated into the estimated mattes to ensure temporal coherence in the final mattes. The proposed NLM framework is shown to reduce temporal jitter when compared to the widely used Laplacian methods using qualitative and quantitative comparisons on a video matting dataset [8], [10].

II. Proposed Approach

A. L1-regularized reconstruction error as alpha matte

The aim of the proposed method is to use reconstruction error as a measure of the α value. The objective of using error reconstruction hinges on the assumption that the foreground and background are locally smooth, akin to propagation-based methods. The idea is therefore, to select a local subset of the known regions for the local smoothness assumption to hold. Following [10], at each pixel, F and B dictionaries are formed by sampling the spatially nearest pixels at a radius of 50 pixels from the definite foreground and background regions, respectively. The feature vector used is the 8-D vector $[\ R\ \ G\ \ B\ \ L\ \ a\ \ b\ \ x\ \ y\ ]^T$ formed by concatenating the RGB and CIELAB color-spaces along with the X-Y coordinates. In order to reduce the sample space, the definite F and B regions are clustered into superpixels using SLIC segmentation [13]. It is to be noted that [10] uses a single dictionary by concatenating the F and B samples together. However, the proposed method requires separate F and B dictionaries in order to determine the reconstruction error with respect to each as explained below.

Given an unknown pixel i, let D_F^i and D_B^i be the foreground and background dictionaries formed by sampling the feature vectors. The sparse codes with respect to each dictionary are determined as

$$\beta_F^i = argmin \left\| v_i - D_F^i \beta_F^i \right\|_2^2 + \lambda \left\| \beta_F^i \right\|_1, \tag{2}$$

$$\beta_B^i = argmin \left\| v_i - D_B^i \beta_B^i \right\|_2^2 + \lambda \left\| \beta_B^i \right\|_1, \tag{3}$$

where v_i is the feature vector at pixel i. The residual errors generated by reconstruction using F and B dictionaries are

$$\xi_F^i = \left\| v_i - D_F^i \beta_F^i \right\|_2, \ \ \xi_B^i = \left\| v_i - D_B^i \beta_B^i \right\|_2. \tag{4}$$

ξ_F^i (ξ_B^i) is the error generated at the unknown pixel i when its feature is reconstructed using foreground (background) dictionary. A high value for ξ_F^i (ξ_B^i) indicates that the current pixel cannot be reconstructed well enough by the F (B) samples. Fig. 2 (b) and (c) visualizes these error maps for a real image. ξ_F^i should ideally be 0 for foreground pixels and gradually increase towards the background pixels. Similarly, ξ_B^i should ideally be 0 for background pixels and gradually increase towards foreground regions. A pixel with a true alpha value of 0.5, i.e. a truly mixed pixel should have comparable reconstruction errors in ξ_F^i and ξ_B^i.

The alpha value can be interpreted as the probability of the pixel belonging to the foreground. ξ_B^i represents the probability of belonging to the foreground, given the known background information, i.e., $\xi_B^i = P(f(i)|\mathbf{D}_B)$. ξ_F^i represents the probability of belonging to the background, given the known foreground information - $\xi_F^i = P(b(i)|\mathbf{D}_F) = 1 - P(f(i)|\mathbf{D}_F)$. Based on the above observation, the alpha value is then estimated as

$$\hat{\alpha}_i = \frac{P(f(i)|\mathbf{D}_B)}{P(f(i)|\mathbf{D}_B) + P(b(i)|\mathbf{D}_F)} = \frac{\xi_B^i}{\xi_B^i + \xi_F^i}. \tag{5}$$

As can be seen from eq. (5), if a pixel truly belongs to the foreground, its foreground reconstruction error ξ_F^i will be a smaller value than the background reconstruction error ξ_B^i, thereby ensuring α is large. The alpha map is shown in Fig. 2(d) and indicates the effectiveness of this simple formulation using sparse reconstruction error.

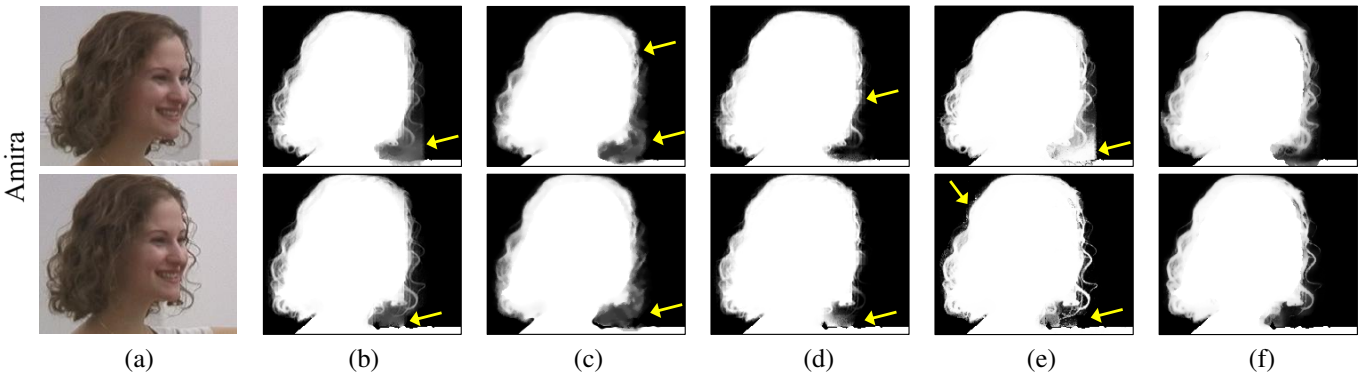

Fig. 4. Qualitative comparison of the video mattes on *Amira* sequence. (a) Input frames, mattes of (b) SC [11], (c) BA [9], (d) EH [8], (e) JO [10] and (f) proposed method. Yellow arrows indicate discontinuities in the matte.

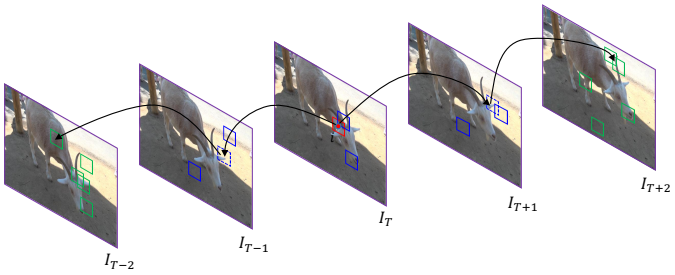

Fig. 3. Illustration of our NLM framework to maintain temporal coherence across the video. For the red patch, AKNN's from immediate frames (blue) are used to find coherent patches from the next frames (green).

B. Patch-based non-local means for temporal coherence

Since the sampling strategy uses a local spatial subset of samples from within the frame, the alpha estimates obtained above lack temporal coherency as the information present in the nearby frames is ignored. Existing methods follow a Laplacian based post-processing step where the inter-pixel correlation is utilized to propagate the matte. The disadvantage inherently lies in its inability to find distant neighbors in space and time. Also, the use of pixel-based matching leads to noise from outliers that get matched incorrectly. To handle this, we propose a patch-based NLM framework that is prevalent in video denoising [14] to maintain the temporal consistency across neighboring frames. NLM [15] was originally introduced to remove noise by averaging pixels in an image weighted by local patch similarities. The high search complexity in finding non-local neighboring patches restricts its use to a local neighborhood alone. Therefore, we apply an approximate K-nearest neighbor patch-matching using coherency sensitive hashing [16] that extends PatchMatch [17] using a hashing scheme where similar patches in the temporal neighborhood are used to propagate the matches to their neighbors.

The framework is illustrated in Fig. 3. For a given image patch $P_T(i)$ (shown in red) centered at pixel i in frame T, approximate K-nearest neighbors (AKNN) in frames $T-1$ and $T+1$ (in blue) are initialized by creating L hash tables based on projection of the patches on Walsh-Hadamard kernels, followed by search for the best candidate patches [16]. The two images are assumed to be coherent, i.e., for every pair of similar patches, their neighbors are also likely to be similar. In the next iteration, the AKNN field for the initial candidates are searched for in their respective temporal neighbors $T-2$ and $T+2$. Under this setup, the temporal neighbors for $P_t(z)$ include a series of AKNN's $\{\mathcal{N}_{T-2}, \mathcal{N}_{T-1}, \mathcal{N}_T, \mathcal{N}_{T+1}, \mathcal{N}_{T+2}\}$. $\mathcal{N}_i = \{P(z_{tj})\}_{j=1}^{K}$ denotes the patches in the t^{th} frame. The non-local means estimate of alpha is [14]

$$\alpha_T(i) = \frac{1}{\Omega} \sum_{t=T-2}^{T+2} \gamma^{|t-T|} \sum_{j=1}^{K} \hat{\alpha}(i_{tj}) \exp\left\{ -\frac{D_w(P(i), P(i_{tj}))}{2\sigma_t^2} \right\}, \tag{6}$$

where $\hat{\alpha}(i)$ is the alpha value of a patch centered at i and Ω is a normalization constant:

$$\Omega = \sum_{t=T-2}^{T+2} \gamma^{|t-T|} \sum_{j=1}^{K} \exp\left\{ -\frac{D_w(P(i), P(i_{tj}))}{2\sigma_t^2} \right\}. \tag{7}$$

$D_w(.,.)$ is a weighted sum of squared difference (SSD) over 2 patches denoted by

$$D_w\left(P(i_1), P(i_2)\right) = \sum_{u \in [-s,s]^2} (P(i_1+u) - P(i_2+u))^2 \exp\left\{ -\frac{\|u\|^2}{2\sigma_p^2} \right\}, \tag{8}$$

where s is the patch width and σ_p is set to $s/2$. γ is set to 0.9 to control the influence of temporal neighbors. Eq. (6) estimates the alpha matte for an entire patch centered at pixel i in frame T. For an alpha patch of width s, all patches whose centers are located within a radius of $\frac{s}{2}$ from i, contain the pixel at i. A simple Gaussian weighted averaging is performed to obtain the final NLM estimate at i as $\alpha_i = \frac{\sum_{j \in \phi} \alpha(i) e^{-\frac{\|i-i_j\|^2}{2\sigma_p^2}}}{\sum_{j \in \phi} e^{-\frac{\|i-i_j\|^2}{2\sigma_p^2}}}$ where ϕ denotes the set of patches that contain the pixel i. The process is repeated for each frame of the video to obtain the final video mattes.

III. Experimental Results

The effectiveness of the proposed method is evaluated on an exclusive video matting dataset used in [8], [10]. It contains sequences covering a wide range of pixel opacity variations

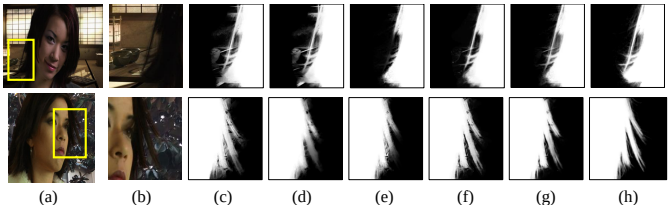

Fig. 5. Qualitative comparison of the video mattes with ground truth on *Face* and *Woman* sequences. (a) Input frames and (b) zoomed region, mattes of (c) SC [11], (d) BA [9], (e) EH [8], (f) JO [10] , (g) proposed method and (h) Ground truth.

and challenges like occlusion and low-contrast. Trimaps are generated on each frame using the method of [18]. λ for sparse coding is set to 0.1. In all experiments, the patch size was set to 8×8. K is set to 5 in eq. (6). The proposed method is evaluated both quantitatively and qualitatively with recent video matting approaches namely, Snapcut (SC) [11], Bai *et al*. (BA) [9], Ehsan *et al*. (EH) [8] and Johnson *et al*. (JO) [10].

A. Qualitative comparison

Fig. 4 and Fig. 5 show the visual comparison of the proposed method against recent video matting methods [8], [10], [9], [11] on *Amira*, *Face* and *Woman* sequences from the dataset. Additional comparisons and video are available at the url[1]. The yellow arrows indicate the regions of discontinuity between consecutive frames. The low contrast between the foreground and the background in *Amira* is a challenging scenario for most matting algorithms. Laplacian-based smoothing used by existing methods produce ambiguity near the boundaries as pixel neighbors tend to be unreliable for accurate propagation. The use of patch-based neighbors in the proposed method enables us to remove such artifacts in the final matte. The sequences in Fig. 5 represent cluttered background which is challenging for most sampling-based algorithms. Our error based matte is accurately able to distinguish between the hair and the background when compared to the ground truth, showing its effectiveness in highly textured regions.

B. Quantitative evaluation

We perform quantitative comparison to evaluate the temporal coherence of the extracted mattes by measuring the difference in alpha values between successive frames as in [19]. The temporal flicker at the i^{th} pixel in frame t is measured as

$$f_i(t) = \frac{|\alpha_i(t+1) - \alpha_i(t)|}{|I_i(t+1) - I_i(t)|}, \tag{9}$$

where $\alpha_i(t)$ and $I_i(t)$ are the alpha and RGB color values at pixel i in frame t.

Table I compares the mean temporal jitter error across 15 sequences in the dataset with recent video matting approaches. As can be seen, the proposed method is able to produce the least temporal jitter across most of the sequences. For the few exceptions, the reconstruction error cannot be trustworthy when the true F and B samples are not present in the dictionary leading to poor initial estimates. Apart from the smooth reconstruction error formulation, the use of a patch-based coherency sensitive hashing is instrumental in the increased performance of the proposed method. [10] uses pixel neighbors in its graph formulation which can be erroneous due to noise. Moreover, the smoothness of alpha is not maintained in the feature vector for sparse coding in [10].

[1]https://goo.gl/Ho5xMN

TABLE I
COMPARISON OF TEMPORAL JITTER ERROR RATES OF DIFFERENT VIDEO MATTING ALGORITHMS AGAINST THE PROPOSED METHOD

Video	SC [11]	BA [9]	EH [8]	JO [10]	Proposed
Face	3.46	2.92	4.26	2.37	**1.49**
Dancer	4.72	4.13	1.48	2.13	**1.46**
Arm	4.43	2.91	2.54	3.52	**1.58**
Woman	4.03	2.72	3.36	2.82	**2.05**
Smoke	3.17	2.96	**1.80**	4.85	2.19
Cat	2.54	4.18	2.45	4.41	**1.40**
Chimp	3.54	4.63	2.90	2.09	**1.81**
Girl	4.55	4.34	2.31	2.12	**1.65**
Whitegoat	3.72	3.85	3.47	2.17	**1.76**
Amira	4.18	4.27	2.72	2.09	**1.72**
Girl2	4.54	4.40	2.18	2.0	**1.86**
Office	3.94	3.64	2.94	**2.05**	2.41
Soccer	4.05	3.24	**2.31**	2.59	2.79
Unicorn	3.23	3.21	2.91	3.34	**2.28**
Dog	4.09	3.62	3.5	2.04	**1.73**

TABLE II
COMPARISON OF RUNNING TIME OF THE PROPOSED METHOD WITH RECENT SAMPLING-BASED APPROACHES

			Total Time (secs)		
Video	Size	No. of Frames	EH [7]	JO [10]	Proposed
Smoke	500x500	90	4491	3798	**3628**
Arm	640x540	49	3260	**1618**	2188
Dancer	480x360	40	5488	2803	**2589**
Face	640x540	78	5378	4955	**4786**
Archaeology	480x405	128	4980	**2524**	2961
Woman	450x400	154	5541	**2912**	3178

Runtime performance: Table II compares the running time of the proposed method with other sampling-based video matting approaches. MATLAB implementations were evaluated on a PC running Intel Xeon 3.2 GHz processor. The proposed method perform comparable to the current approaches without compromising on the quality of the matte.

IV. CONCLUSION

We present a novel video matting framework that treats the matting problem from the perspective of reconstruction error of a feature. Foreground and background dictionaries, whose bases are used to reconstruct an unknown feature vector with L1-regularization are used to measure the error towards F and B respectively. A NLM framework is also proposed that is integrated across multiple frames to ensure temporal coherence in the video mattes. Experimental evaluations demonstrate that the proposed method has advantages over current matting methods that use a Laplacian based smoothing.

REFERENCES

[1] J. Wang and M. F. Cohen, "Optimized color sampling for robust matting," in *Proc. IEEE CVPR*, 2007, pp. 1–8.

[2] K. He, C. Rhemann, C. Rother, X. Tang, and J. Sun, "A global sampling method for alpha matting," in *Proc. IEEE CVPR*, 2011, pp. 2049–2056.

[3] E. Shahrian, D. Rajan, B. Price, and S. Cohen, "Improving image matting using comprehensive sampling sets," in *Proc. IEEE CVPR*, 2013, pp. 636–643.

[4] J. Johnson, D. Rajan, and H. Cholakkal, "Sparse codes as alpha matte," in *Proc. BMVC*, 2014.

[5] A. Levin, D. Lischinski, and Y. Weiss, "A closed-form solution to natural image matting," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 2, pp. 228–242, 2008.

[6] Q. Chen, D. Li, and C.-K. Tang, "Knn matting," in *Proc. IEEE CVPR*, 2012, pp. 869–876.

[7] D. Cho, Y.-W. Tai, and I. Kweon, "Natural image matting using deep convolutional neural networks," in *Proc. ECCV*, 2016, pp. 626–643.

[8] E. Shahrian, B. Price, S. Cohen, and D. Rajan, "Temporally coherent and spatially accurate video matting," *Computer Graphics Forum*, vol. 33, no. 2, pp. 381–390, 2014.

[9] X. Bai, J. Wang, and D. Simons, "Towards temporally-coherent video matting," in *Computer Vision/Computer Graphics Collaboration Techniques*. Springer, 2011, pp. 63–74.

[10] J. Johnson, E. S. Varnousfaderani, H. Cholakkal, and D. Rajan, "Sparse coding for alpha matting," *IEEE Trans. Image Process.*, vol. 25, no. 7, pp. 3032–3043, July 2016.

[11] X. Bai, J. Wang, D. Simons, and G. Sapiro, "Video snapcut: robust video object cutout using localized classifiers," in *ACM Trans. Graph.*, vol. 28, no. 3. ACM, 2009, p. 70.

[12] D. Li, Q. Chen, and C.-K. Tang, "Motion-aware knn laplacian for video matting," in *Proc. IEEE ICCV*, 2013, pp. 3599–3606.

[13] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Susstrunk, "Slic superpixels compared to state-of-the-art superpixel methods," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 11, pp. 2274–2282, 2012.

[14] C. Liu and W. T. Freeman, "A high-quality video denoising algorithm based on reliable motion estimation," in *Proc. ECCV*, 2010, pp. 706–719.

[15] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *Proc. IEEE CVPR*, vol. 2, 2005, pp. 60–65.

[16] S. Korman and S. Avidan, "Coherency sensitive hashing," in *Proc. IEEE ICCV*, 2011, pp. 1607–1614.

[17] C. Barnes, E. Shechtman, D. B. Goldman, and A. Finkelstein, "The generalized patchmatch correspondence algorithm," in *Proc. ECCV*, 2010, pp. 29–43.

[18] J. Johnson, D. Rajan, and H. Cholakkal, "Temporal trimap propagation using motion-assisted shape blending," in *Proc. IEEE VCIP*, 2015, pp. 1–4.

[19] S.-Y. Lee, J.-C. Yoon, and I.-K. Lee, "Temporally coherent video matting," *Graph. Models.*, vol. 72, no. 3, pp. 25–33, 2010.